

Interpolation in Equilibrium Logic and Answer Set Programming: the Propositional Case

Dov Gabbay¹, David Pearce^{2*} and Agustín Valverde^{3**}

¹ King's College London

dov.gabbay@kcl.ac.uk

² AI Dept, Universidad Politécnica de Madrid, Spain.

david.pearce@upm.es

³ Dept of Applied Mathematics, Universidad de Málaga, Spain.

a.valverde@ctima.uma.es

Abstract. Interpolation is an important property of classical and many non classical logics that has been shown to have interesting applications in computer science and AI. Here we study the Interpolation Property for the propositional version of the non-monotonic system of *equilibrium logic*, establishing weaker or stronger forms of interpolation depending on the precise interpretation of the inference relation. These results also yield a form of interpolation for ground logic programs under the answer sets semantics. For disjunctive logic programs we also study the property of *uniform* interpolation that is closely related to the concept of variable forgetting.

1 Introduction

The Interpolation Property plays an important role in logical systems, both classical and non-classical. Its importance in computer science has also become recognised lately. The Interpolation Property has been applied in various areas of computer science, notably in software specification [4] and in the construction of formal ontologies [12]. In both cases it is relevant to modularity issues, for instance in [12] it plays a key role in the study of the modular decomposition of ontologies. However to date interpolation has received less attention in systems of nonmonotonic reasoning and logic programming, despite their importance in AI and computer science. In this paper we study the interpolation property for the system of nonmonotonic reasoning known as *equilibrium logic*.⁴ Since this in turn forms a logical foundation for stable model reasoning and Answer Set Programming (ASP), our results transfer easily to the sphere of ASP.

Let us start with some notation and terminology. Let us assume the syntax of propositional logic with formulas denoted by lower case Greek letters and let \vdash be a monotonic inference relation. If $\alpha \vdash \beta$, an *interpolant for* (α, β) is a formula γ such that

$$\alpha \vdash \gamma \ \& \ \gamma \vdash \beta \tag{1}$$

* Partially supported by CICyT projects TIC-2003-9001-C02 and TIN2006-15455-CO3.

** Partially supported by CICyT project TIC-2003-9001-C01, TIN2006-15455-CO1 and Junta de Andalucía project TIC-115

⁴ We are very grateful to Karl Schlechta for valuable discussions on the interpolation concept for non-monotonic logics.

where γ contains only variables that belong to both α and β . A logic L with inference relation \vdash_L is said to have the *interpolation property* if an interpolant exists for every pair of formulas (α, β) such that $\alpha \vdash_L \beta$. As is well-known, classical logic as well as many non-classical logics possess interpolation [10].

Suppose now we deal with a non-monotonic logical system with an inference relation \vdash . To express the idea that a formula is an interpolant one approach is simply to replace \vdash by \vdash in (1). We may call this form (\vdash, \vdash) -interpolation. Although this is a legitimate form of interpolation, in practice we shall find it useful to study another approach suggested in [10]. Suppose that our non-monotonic inference can be defined in terms of minimal models in some monotonic logical system, say that the relation \vdash is captured by means of minimal models in a logic L with inference relation \vdash_L ⁵. Assume $\alpha \vdash \beta$; then as an interpolant for (α, β) we look for a formula γ such that

$$\alpha \vdash \gamma \ \& \ \gamma \vdash_L \beta \quad (2)$$

where all variables of γ occur in both α and β . Since \vdash is to be defined via a subclass of minimal L -models, we already suppose that $\models_L \subseteq \vdash$. Moreover we should require too that L is a well-behaved sublogic in the sense that L -equivalent formulas have the same \vdash -consequences and that formulas L -derivable from \vdash -consequences are themselves \vdash -consequences (so eg from (2) we can derive $\alpha \vdash \beta$). In non-monotonic reasoning these last two properties are known as left and right absorption, respectively [15]. Given these conditions, it follows at once from (2) that any formula in the language of γ that is L -equivalent to γ will also be an interpolant for (α, β) . Likewise if γ is an interpolant for (α, β) and $\beta \vdash_L \delta$ then $\alpha \vdash \delta$ and γ is an interpolant for (α, δ) . When (2) holds we call γ a (\vdash, \vdash_L) -interpolant.

Now, to find a (\vdash, \vdash_L) -interpolant for (α, β) , we can proceed as follows. We look for an L -formula α' say, that precisely L -defines the minimal models of α . Since $\alpha \vdash \beta$ it follows that $\alpha' \models_L \beta$ and, assuming completeness, $\alpha' \vdash_L \beta$. Now, if L has the interpolation property as defined earlier, we apply this theorem to obtain or infer the existence of an L -interpolant γ in the sense of (1) for (α', β) . Hence (2) follows.

1.1 Scope of the paper

In this paper we study interpolation for the propositional version of equilibrium logic, based on the non-classical, monotonic logic of *here-and-there*, **HT**. We introduce two variants of equilibrium inference, denoted by \vdash_{cw} and \vdash_{ow} respectively. This notation reflects the idea that one form of inference is closer in spirit to closed-world reasoning while the other resembles more a type of open-world reasoning. While (\vdash, \vdash_L) -interpolation holds for \vdash_{ow} (setting $L = \mathbf{HT}$), only the weaker form of (\vdash, \vdash) -interpolation holds for \vdash_{cw} . However in both cases we apply the general method described above, using definable classes of minimal **HT**-models.

The restriction to propositional equilibrium logic is quite sufficient for considering interpolation in ASP, for the case of finite, ground logic programs (of any syntactic kind). In the case of ASP, the most natural associated form of inference would seem

⁵ Assume by completeness that this coincides with L -consequence, \models_L .

to be \vdash_{cw} , satisfying (\vdash, \vdash) -interpolation. As a final topic we consider the extent to which a stronger form of *uniform* interpolation holds for disjunctive programs under a restricted query language. Here we make use of recent results by Eiter and Wang [8] on variable forgetting in ASP.

For reasons of space we do not consider here the full, first-order version of equilibrium logic that can serve as a foundation for non-ground answer set programs. This is done in the full, extended version of the paper that is currently in preparation.

2 Logical Preliminaries

We work with standard propositional languages, $\mathcal{L}, \mathcal{L}'$, etc based on a sets V, V' , of propositional variables.⁶ Formulas are built-up in the usual way using the logical constants $\wedge, \vee, \rightarrow, \neg$, standing respectively for conjunction, disjunction, implication and negation. If φ is a propositional formula, we denote by $V(\varphi)$ the set of propositional variables appearing in φ .

As usual the symbols \vdash and \models , possibly with subscripts, are used to denote logical inference and consequence relations, respectively. A logic L is said to be *monotonic* if its inference relation \vdash_L satisfies the monotonicity property:

$$\Pi \vdash_L \varphi \ \& \ \Pi \subseteq \Pi' \rightarrow \Pi' \vdash_L \varphi \quad (3)$$

To distinguish non-monotonic from monotonic inference relations, we use \vdash to symbolise the former. In most cases a non-monotonic logic can be understood in terms of an inference relation that extends a suitable monotonic logic. When this extension is well-behaved we say that the monotonic logic forms a *deductive base* for it. This can be made precise as follows.

Definition 1. *Let \vdash be any nonmonotonic inference relation. We say that a logic L with monotonic inference relation \vdash_L is a deductive base for \vdash iff (i) $\vdash_L \subseteq \vdash$; (ii) If $\Pi_1 \equiv_L \Pi_2$ then $\Pi_1 \approx \Pi_2$; (iii) If $\Pi \vdash \varphi$ and $\varphi \vdash_L \psi$, then $\Pi \vdash \psi$.*

Here \equiv_L denotes ordinary logical equivalence in L , while \approx denotes non-monotonic equivalence, ie $\Pi_1 \approx \Pi_2$ means that Π_1 and Π_2 have the same non-monotonic consequences. Furthermore, we say that a deductive base is *strong* if it satisfies the additional condition:

$$\Pi_1 \not\equiv_L \Pi_2 \rightarrow \text{there exists } \Gamma \text{ such that } \Pi_1 \cup \Gamma \not\approx \Pi_2 \cup \Gamma.$$

In terms of nonmonotonic consequence operations, (ii) and (iii) correspond to conditions known as left absorption and right absorption respectively, see [15].⁷

⁶ Many logic texts work with a fixed, countable set of propositional variables. However here we find it useful to distinguish different languages $\mathcal{L}, \mathcal{L}'$, and variable sets V, V' , etc. For one thing, these languages may arise by grounding or instantiating finite, first-order theories, as occurs in ASP, hence it may be important to distinguish the different languages that may result from this process. Secondly, in a non-monotonic context, as we shall see, different kinds of inference relations may arise according to the way in which language extensions are handled. Thirdly, our definitions will be easily extended to the first-order case.

⁷ In the terminology of [7] we therefore require of \vdash_L, \vdash that they form a *fully absorbing inferential frame*.

We now turn to the interpolation property.

Definition 2. A logic L with inference relation \vdash_L is said to have the interpolation property if whenever

$$\vdash_L \varphi \rightarrow \psi$$

there exists a sentence ξ (the interpolant) such that $V(\xi) \subseteq V(\varphi) \cap V(\psi)$ and

$$\vdash_L \varphi \rightarrow \xi \quad \text{and} \quad \vdash_L \xi \rightarrow \psi$$

As explained in the introduction, for non-monotonic logics we can consider two forms of interpolation, one weaker one stronger. The stronger form makes use of an underlying monotonic logic.

Definition 3. Suppose that $\alpha \vdash \beta$. A (\vdash, \vdash_L) interpolant for (α, β) is a formula γ such that $V(\gamma) \subseteq V(\alpha) \cap V(\beta)$ and

$$\alpha \vdash \gamma \quad \text{and} \quad \gamma \vdash_L \beta \tag{4}$$

where L is a deductive base for \vdash . A non-monotonic logic with inference relation \vdash is said to have the (\vdash, \vdash_L) interpolation property if for a suitable deductive base logic L a (\vdash, \vdash_L) interpolant exists for every pair of formulas (α, β) such that $\alpha \vdash \beta$.

The requirement that L form a deductive base ensures that some desirable properties of interpolation are met.

Proposition 1. Let γ be a (\vdash, \vdash_L) interpolant for (α, β) .

1. For any ψ such that $\psi \equiv_L \gamma$, ψ is a (\vdash, \vdash_L) interpolant for (α, β) .
2. For any α' such that $\alpha \equiv_L \alpha'$, and any β' such that $\beta \vdash_L \beta'$, γ is a (\vdash, \vdash_L) interpolant for (α', β') .

The property of deductive base also guarantees that the (\vdash, \vdash_L) relation is transitive in the sense that if (4) holds for any α, β, γ , then also $\alpha \vdash \beta$. This last property will not necessarily hold for the second, weaker form of interpolation that we call (\vdash, \vdash) interpolation.

Definition 4. Suppose that $\alpha \vdash \beta$. A (\vdash, \vdash) interpolant for (α, β) is a formula γ such that $V(\gamma) \subseteq V(\alpha) \cap V(\beta)$ and

$$\alpha \vdash \gamma \quad \text{and} \quad \gamma \vdash \beta \tag{5}$$

Analogous to the previous case, we say that a non-monotonic logic with inference relation \vdash has the (\vdash, \vdash) interpolation property if a (\vdash, \vdash) interpolant exists for every pair of formulas (α, β) such that $\alpha \vdash \beta$. Notice that (\vdash, \vdash_L) is the stronger form of interpolation because if a logic has (\vdash, \vdash_L) interpolation it must also have (\vdash, \vdash) interpolation, again as a consequence of the deductive base requirement (first clause).

Evidently the properties expressed in Proposition 1 are not directly applicable to the second form of interpolation that does not refer to any underlying base logic. Nevertheless an important feature of the interpolation properties we shall establish below is that we can formulate and prove analogous properties even for (\vdash, \vdash) interpolation.

We can also consider restricted variants of interpolation when the property holds for certain types of formulas, in other words, when there is an interpolant for (α, β) whenever α and β belong to specific syntactic classes. Later on we shall consider both kinds of restrictions, where α belongs to a specific class or alternatively when β does.

2.1 Review of the Logic of Here-and-There

Equilibrium logic is based on the nonclassical logic of here-and-there, which we denote by **HT** in the propositional case. The axioms and rules of inference for **HT** are those of intuitionistic logic together with the axiom schema

$$\alpha \vee (\neg\beta \vee (\alpha \rightarrow \beta)).$$

The model theory of **HT** is based on the usual Kripke semantics for intuitionistic logic [21], but it is complete for Kripke frames with just two worlds h (here) and t (there) such that $h \leq t$. We can therefore represent a Kripke model for **HT** as a triple $\langle \{h, t\}, \leq, v \rangle$, where v is a (truth) valuation. Alternatively, we can speak of an **HT-interpretation** as an ordered pair $\mathcal{M} = \langle H, T \rangle$ of sets of atoms such that $H \subseteq T$; the elements of H are the atoms true *here* and the elements of T are the atoms true *there*. The truth of a formula φ in a world $\mathcal{M}, w \models \varphi$ is defined recursively via the usual rules for conjunction, disjunction, implication and negation in intuitionistic logic. A formula φ is true in $\mathcal{M} = \langle H, T \rangle$ in symbols $\mathcal{M} \models \varphi$, if it is true at each world in \mathcal{M} ; in this case we say that \mathcal{M} is an **HT-model** of φ . A formula φ is said to be *valid* in **HT**, in symbols $\models \varphi$, if it is true in all **HT-interpretations**. Logical consequence for **HT** is understood as follows: φ is said to be an **HT** consequence of a theory Π , written $\Pi \models \varphi$, iff for all models \mathcal{M} and any world $w \in \mathcal{M}$, $\mathcal{M}, w \models \Pi$ implies $\mathcal{M}, w \models \varphi$. Equivalently this can be expressed by saying that φ is true in all **HT-models** of Π .

More exactly, we might write $\mathbf{HT}(\mathcal{L})$ to refer to the language considered in the logic. However, as we see below, the logic is in fact independent from the language.

Let \mathcal{L} be a proper sublanguage of \mathcal{L}' , ie $\mathcal{L} \subset \mathcal{L}'$; for any **HT**(\mathcal{L}')-interpretation $\mathcal{M} = \langle H, T \rangle$ we denote by $\mathcal{M}|_{\mathcal{L}}$ the **HT**(\mathcal{L})-interpretation formed by omitting the interpretation of all atoms in $\mathcal{L}' \setminus \mathcal{L}$ and we call this the *reduct* of \mathcal{M} to \mathcal{L} .

Proposition 2. *Suppose that $\mathcal{L}' \supset \mathcal{L}$, Π is a theory in \mathcal{L} and \mathcal{M} is an **HT**(\mathcal{L}')-model of Π . Then $\mathcal{M}|_{\mathcal{L}}$ is a **HT**(\mathcal{L})-model of Π .*

Proposition 3. *Suppose that $\mathcal{L}' \supset \mathcal{L}$ and $\varphi \in \mathcal{L}$. Then φ is valid (resp. satisfiable) in **HT**(\mathcal{L}) if and only if is valid (resp. satisfiable) in **HT**(\mathcal{L}').*

A key property of here-and-there is that, as Maksimova [16] showed, it is one of just seven super-intuitionistic logics with Interpolation.

Proposition 4 ([16]). *The logic **HT** possesses the Interpolation Property.*

2.2 Equilibrium Logic

Equilibrium logic is based on certain kinds of minimal models in **HT**.

Definition 5. *Among here-and-there interpretations we define the order \trianglelefteq as follows: $\langle H, T \rangle \trianglelefteq \langle H', T' \rangle$ if $T = T'$ and $H \subseteq H'$. If $H \subset H'$ we write $\langle H, T \rangle \triangleleft \langle H', T' \rangle$*

Definition 6 (Equilibrium model). *Let Π be a theory and $\mathcal{M} = \langle H, T \rangle$ a model of Π .*

1. \mathcal{M} is said to be *total* if $H = T$.

2. \mathcal{M} is said to be an equilibrium model of Π if it is minimal under \leq among models of Π , and it is total.

In other words, equilibrium models are total models for which there is no ‘smaller’ non-total model. Evidently a total **HT**-model of a theory Π can be equivalently regarded as a classical model of Π ; and in what follows we make tacit use of this equivalence. A theory is said to be *consistent* if it has an **HT**-model and *coherent* if it has an equilibrium model.

We define a preliminary notion of equilibrium entailment as follows. It essentially agrees with standard versions of equilibrium logic, as eg in [18].

Definition 7. *The relation \vdash , called equilibrium entailment, is defined as follows. Let Π be a set of formulas.*

1. *If Π is non-empty and has equilibrium models, then $\Pi \vdash \varphi$ if every equilibrium model of Π is a model of φ in **HT**.*
2. *If either Π is empty or has no equilibrium models, then $\Pi \vdash \varphi$ if $\Pi \vdash \varphi$.*

A few words may help to explain the concept of equilibrium entailment. First, we define the basic notion of entailment as truth in every intended (equilibrium) model. In nonmonotonic reasoning this is a common approach and sometimes called a *skeptical* or *cautious* notion of entailment or inference; its counterpart *brave* reasoning being defined via truth in some intended model. Since equilibrium logic is intended to provide a logical foundation for the answer set semantics of logic programs, the cautious variant of entailment is the natural one to choose: the standard consequence relation associated with answer sets is given by truth in all answer sets of a program. Note however that in ASP as a programming paradigm each answer set may correspond to a particular solution of the problem being modelled and is therefore of interest in its own right.

Secondly, it is useful to have a nonmonotonic consequence or entailment relation that is non-trivially defined for all consistent theories. As is easily seen, however, not all such theories possess equilibrium models. For such cases it is natural to use monotonic consequence as the entailment relation. In particular, **HT** is a maximal logic with the property that logically equivalent theories have the same equilibrium models. Evidently situation 2 in previous definition also handles correctly the cases that Π is empty or inconsistent.

Despite these qualifications, there remains an ambiguity in the concept of equilibrium entailment that we now need to settle. Suppose that $\mathcal{L}' \supset \mathcal{L}$, Π is a theory in \mathcal{L} and φ is a sentence in \mathcal{L}' . How should we understand the expression ‘ $\Pi \vdash \varphi$ ’?

Evidently, if we fix a language in advance, say as the language \mathcal{L}' , then we can simply consider the equilibrium models of Π in \mathcal{L}' . But if Π represents a knowledge base or a logic program, for instance, we may also take the view that $V(\Pi)$ is the appropriate language to work with. In that case, the query φ is as such not fully interpreted as it contains some variables not in $V(\Pi)$.

For any language \mathcal{L} and \mathcal{L} -theory Π , let $E_{\mathcal{L}}(\Pi)$ be the collection of all equilibrium models of Π in **HT**(\mathcal{L}). Now consider the following two variants of entailment.

Definition 8 (Equilibrium entailment). *Assume Π is non-empty and has equilibrium models, then:*

1. Let us say that $\Pi \vdash_{cw} \varphi$ if and only if $\mathcal{M} \models \varphi$ for each $\mathcal{M} \in E_{\mathcal{L}'}(\Pi)$, where \mathcal{L}' is the language over $V(\Pi \cup \{\varphi\})$.
2. Let us say that $\Pi \vdash_{ow} \varphi$ if and only if $\mathcal{M} \models \varphi$ for each $\mathcal{M} \in E_{\mathcal{L}}(\Pi) \upharpoonright^{\mathcal{L}'}$, where \mathcal{L}' is the language over $V(\varphi)$ and $E_{\mathcal{L}}(\Pi) \upharpoonright^{\mathcal{L}'}$ denotes the collection of all expansions of elements of $E_{\mathcal{L}}(\Pi)$ to models in $\mathcal{L} \cup \mathcal{L}'$, ie where the vocabulary of $\mathcal{L}' \setminus \mathcal{L}$ is interpreted arbitrarily.

Obviously, if either Π is empty or has no equilibrium models, then $\Pi \vdash_{cw} \varphi$ iff $\Pi \vdash_{ow} \varphi$ iff $\Pi \vdash \varphi$.

A simple example will illustrate the difference between \vdash_{cw} and \vdash_{ow} . Let ψ be an \mathcal{L} -formula and let q be a variable not in \mathcal{L} and let \mathcal{L}' be the language $\mathcal{L} \cup \{q\}$. By the first method we have $\psi \vdash_{cw} \psi \wedge (q \vee \neg q)$. In fact we have the stronger entailment $\psi \vdash_{cw} \psi \wedge \neg q$. The reason is that when we form the equilibrium models of ψ in \mathcal{L}' , q will be false in each as an effect of taking minimal models. On the other hand, if we expand equilibrium models of ψ in $\mathbf{HT}(\mathcal{L})$ to $\mathbf{HT}(\mathcal{L}')$, the new variable q receives an arbitrary interpretation in $\mathbf{HT}(\mathcal{L}')$. Since this logic is 3-valued we do not obtain $\Pi \vdash_{ow} q \vee \neg q$.

For standard, monotonic logics, there is no difference between these two forms of entailment. If in Definition 8 we replace everywhere equilibrium model by simply model (in \mathbf{HT}), variants (i) and (ii) give the same result.

In the context of logic programming and deductive databases the more orthodox view is that reasoning is based on a *closed world assumption* (CWA). Accordingly a ground atomic query like $q(a)?$, where the predicate q or the name a do not belong to the language of the program or database would simply be assigned the value *false*. This is also the case with the first kind of equilibrium entailment and we use the label \vdash_{cw} since this variant appears closer to a closed world form of reasoning. On the other hand, there may be legitimate cases where we do not want to apply the CWA and where unknown values should be assigned to an atom that is not expressed in the theory language. Then the second form of entailment, \vdash_{ow} , nearer to open world reasoning, may then be more appropriate. For present purposes, however, the suffices '*cw*' and '*ow*' should be thought of merely as mnemonic labels. A more thorough analysis of closed world versus open world reasoning in this context would probably lead us to consider assumptions such as *unique names assumption* or *standard names assumption* and is outside the scope of this paper.

3 Interpolation in Propositional Equilibrium Logic

In this section we deal with interpolation in propositional equilibrium logic. It is clear that by its semantic construction propositional equilibrium logic has \mathbf{HT} as a deductive base. This base is actually maximal.

Proposition 5. *\mathbf{HT} is a strong and maximal deductive base for (propositional) equilibrium entailment.*

The first property is precisely the strong equivalence theorem of [13]. Maximality follows from the fact that any logic strictly stronger than \mathbf{HT} would have to contain classical logic which is easily seen not to be a deductive base, eg violating condition (ii) of Definition 1. We have:

Lemma 1. *Let α be a coherent **HT**-formula and $E(\alpha)$ its set of equilibrium models. Then there is formula α' of **HT** in $V(\alpha)$ that defines $E(\alpha)$ in the sense that $\mathcal{M} \in E(\alpha)$ iff $\mathcal{M} \models \alpha'$.*

Proof. Suppose that α is coherent. and let

$$\mathcal{M}_1 = \langle T_1, T_1 \rangle, \mathcal{M}_2 = \langle T_2, T_2 \rangle, \dots, \mathcal{M}_n = \langle T_n, T_n \rangle$$

be an enumeration of its equilibrium models. We show how to define $E(\alpha)$. Suppose each T_i , has k_i elements and denote them by $A_1^i, \dots, A_j^i, \dots, A_{k_i}^i$. Let \bar{T}_i be the complement of T_i ; then we can list its members as $A_{k_i+1}^i, \dots, A_l^i, \dots, A_{|V(\alpha)|}^i$. Set

$$\delta^i = \bigwedge_{j=1, \dots, k_i} A_j^i \wedge \neg \left(\bigvee_{l=k_i+1, \dots, |V(\alpha)|} A_l^i \right) \quad (6)$$

Now set

$$\alpha' = \bigvee_{i=1, \dots, n} \delta^i \quad (7)$$

We claim that $\mathcal{M} \models \alpha'$ if and only if $\mathcal{M} = \mathcal{M}_i$ for some $i = 1, \dots, n$, ie the models of α' are precisely $\mathcal{M}_1, \dots, \mathcal{M}_n$. To verify this claim, note that each $\mathcal{M}_i \models \delta^i$ and so $\mathcal{M}_i \models \alpha'$. Conversely, suppose that $\mathcal{M} \models \alpha'$. From the semantics of **HT** it is clear that $\mathcal{M} \models \varphi \vee \psi$ iff $\mathcal{M} \models \varphi$ or $\mathcal{M} \models \psi$, so in particular $\mathcal{M} \models \alpha'$ implies $\mathcal{M} \models \delta^i$ for some $i = 1, \dots, n$. However, each δ^i defines a complete theory whose models are total. It follows that if $\mathcal{M} \models \delta^i$, then $\mathcal{M} = \mathcal{M}_i$. This establishes the claim. \square

Although we shall now demonstrate interpolation in the (\vdash, \vdash) form for the relation \vdash_{cw} , we actually establish a stronger result. One consequence of this is that if we are concerned with \vdash_{ow} entailment then the (\vdash, \vdash) form of interpolation actually holds.

Proposition 6 ((\vdash, \vdash) -Interpolation). *Let α, β be formulas and set $V = V(\alpha) \cup V(\beta)$ and $V' = V(\beta) \setminus V(\alpha)$ and suppose that B_1, \dots, B_n is an enumeration of V' . If $\alpha \vdash_{cw} \beta$, there is a formula γ such that $V(\gamma) \subseteq V(\alpha) \cap V(\beta)$, $\alpha \vdash \gamma$, and $\gamma \wedge \neg B_1 \wedge \dots \wedge \neg B_n \models \beta$. Hence in particular $\gamma \vdash_{cw} \beta$.*

Proof. Let α, β and V, V' be as in the statement of the proposition, and suppose that $\alpha \vdash_{cw} \beta$. Then β holds in all equilibrium models of α in the language V . Case (i): suppose that α is coherent and form its set of equilibrium models, $E_V(\alpha)$. By the equilibrium construction it is easy to see that in each model $\mathcal{M} \in E_V(\alpha)$ each atom B_i is false, for $i = 1, n$. Construct the formulas δ_i and the formula α' exactly as in the proof of Lemma 1. Now consider the formula $(\neg B_1 \wedge \dots \wedge \neg B_n) \wedge \alpha'$. Clearly this formula defines the set of equilibrium models of α in **HT**(V). Consequently, $(\neg B_1 \wedge \dots \wedge \neg B_n) \wedge \alpha' \models \beta$ and so $\alpha' \vdash (\neg B_1 \wedge \dots \wedge \neg B_n) \rightarrow \beta$. We can now apply the interpolation theorem for **HT** to infer that there is a formula γ such that $\alpha' \vdash \gamma$ and $\gamma \vdash (\neg B_1 \wedge \dots \wedge \neg B_n) \rightarrow \beta$, where $V(\gamma) \subseteq V(\alpha') \cap V(\beta)$ and hence $V(\gamma) \subseteq V(\alpha) \cap V(\beta)$. Since **HT** is a deductive base, we conclude that

$$\alpha \vdash \gamma \ \& \ \gamma \wedge \neg B_1 \wedge \dots \wedge \neg B_n \vdash \beta.$$

Now, since $V(\gamma) \subseteq V(\alpha) \cap V(\beta)$, $B_i \notin V(\gamma)$ for $i = 1, \dots, n$. It follows that in $\mathbf{HT}(V(\beta))$, each B_i is false in every equilibrium model of γ . So each such model \mathcal{M} satisfies $(\neg B_1 \wedge \dots \wedge \neg B_n)$.⁸ Since each also satisfies β , we have $\gamma \vdash_{cw} \beta$.

Case (ii). If α has no equilibrium models then the hypothesis is that $\alpha \vdash \beta$. In that case we simply choose an interpolant γ for (α, β) . \square

Corollary 1 ((\vdash, \vdash)-Interpolation). *Let α, β be formulas such that $\alpha \vdash_{cw} \beta$ and $V(\beta) \subseteq V(\alpha)$. There is a formula γ such that $V(\gamma) \subseteq V(\alpha) \cap V(\beta)$ and $\alpha \vdash_{cw} \gamma$ and $\gamma \vdash \beta$.*

Proof. Immediate from Proposition 6 by the fact that $V(\beta) \setminus V(\alpha) = \emptyset$. \square

Proposition 7 ((\vdash, \vdash)-Interpolation). *Let α, β be formulas and set $V = V(\alpha) \cup V(\beta)$ and $V' = V(\beta) \setminus V(\alpha)$. If $\alpha \vdash_{ow} \beta$, there is a formula γ such that $V(\gamma) \subseteq V(\alpha) \cap V(\beta)$, $\alpha \vdash_{ow} \gamma$, and $\gamma \vdash \beta$.*

Proof. Let α, β and V, V' be as in the statement of the proposition and suppose that $\alpha \vdash_{ow} \beta$. Then β holds in all expansions of elements of $E_{V(\alpha)}(\alpha)$ to the language V . Case (i): suppose that α is coherent and consider $E_{V(\alpha)}(\alpha)$. Again construct the formulas δ_i and the formula α' exactly as in the proof of Lemma 1. Now consider α' which defines the set $E_{V(\alpha)}(\alpha)$. Then β holds in all expansions of models of α' to V . Hence $\alpha' \models \beta$ and therefore $\alpha' \vdash \beta$. We can now apply the interpolation theorem for \mathbf{HT} to infer that there is a formula γ such that $\alpha' \vdash \gamma$ and $\gamma \vdash \beta$, where $V(\gamma) \subseteq V(\alpha') \cap V(\beta)$ and hence $V(\gamma) \subseteq V(\alpha) \cap V(\beta)$. Since $\alpha \vdash_{ow} \alpha'$ and \mathbf{HT} is a deductive base we conclude that

$$\alpha \vdash_{ow} \gamma \ \& \ \gamma \vdash \beta.$$

Case (ii). If α has no equilibrium models, choose γ as an interpolant for (α, β) . \square

4 Interpolation in Answer Set Semantics

Answer set programming (ASP) has become an established form of declarative, logic-based programming and its basic ideas are now well-known. For a textbook treatment the reader is referred to [3]. As is also well-known, the origins of ASP lie in the stable model and answer set semantics for logic programs introduced in [11]. This semantics made use of a fixpoint condition involving a certain ‘reduct’ operator. Subsequent extensions of the concept to cover more general kinds of rules also relied on a reduct operator of similar sort [14,9]. For the original definitions, the reader is referred to the various papers cited.

The correspondence between answer set semantics and equilibrium logic is also well-established and has been discussed in many publications, beginning with [17] which first showed how the answer sets of disjunctive programs can be regarded as equilibrium models. For our purposes it will suffice to recall just the main features of the correspondence with equilibrium logic.

⁸ Notice that in this case adding to γ the sentence $(\neg B_1 \wedge \dots \wedge \neg B_n)$ does not change its set of equilibrium models.

We recall the notion of ground, disjunctive logic program (without strong negation) whose answer sets are simply collections of atoms. These programs consist of sets of ground rules of the form

$$K_1 \vee \dots \vee K_k \leftarrow L_1, \dots, L_m, \text{not} L_{m+1}, \dots, \text{not} L_n \quad (8)$$

where the L_i and K_j are atoms. The ‘translation’ from the syntax of programs to **HT** propositional formulas is the trivial one, viz. (8) corresponds to the **HT** sentence

$$L_1 \wedge \dots \wedge L_m \wedge \neg L_{m+1} \wedge \dots \wedge \neg L_n \rightarrow K_1 \vee \dots \vee K_k \quad (9)$$

Under this translation the correspondence between the answer sets and the equilibrium models of ground disjunctive programs is also the direct one:

Proposition 8. *Let Π be a disjunctive logic program. Then $\langle T, T \rangle$ is an equilibrium model of Π if and only if T is an answer set of Π .*

This was first shown in [17], but the basic equivalence was later shown to hold for more general classes of programs in [13]. Indeed it can also be extended to embrace the very general definition of answer set for propositional theories, given by Ferraris [9].

In ASP the main emphasis is on finding answer sets and this is what most answer set solvers compute. Less attention is placed on implementing a non-monotonic inference relation.⁹ However there is a standard, skeptical concept of inference or entailment associated with answer set semantics. This notion of entailment or consequence for programs under the answer set semantics is that a query Q is entailed by a program Π if Q is true in all answer sets of Π , see eg [2]. Let us denote this entailment or consequence relation by \vdash_{AS} . Evidently atoms are true in an answer set if and only if they belong to it. Conjunctions and disjunctions are handled in the obvious way (eg [14,2]). Sometimes, as in [2], queries of the form *not* a , or in logical notation $\neg a$, are not explicitly dealt with. However it seems to be in keeping with the semantics to regard a formula of form $\neg a$ to be true in an answer set if and only if a is not true. Another way to express this would be to say that an answer set satisfies $\neg \alpha$ if it does not violate the constraint $\{\leftarrow \alpha\}$, where constraint violation is understood as in [14].¹⁰ In this way we would say that $\Pi \vdash_{AS} \neg \alpha$ if no answer set of Π contains α .

We can therefore transfer interpolation properties from equilibrium logic to answer set semantics and ASP. It remains to consider whether \vdash_{AS} is best identified with the closed world version of inference, \vdash_{cw} , or the more open world version, \vdash_{ow} . Again, since ASP solvers do not implement inference engines, the difference is really a theoretical one. In traditional logic programming, however, a query that does not belong to the language of the program is usually answered *false*. It also seems quite natural in an ASP context that, given a program Π and a query Q , one should consider the stable models of Π in the language $V(\Pi) \cup V(Q)$ even if this is a proper extension of the language of Π .¹¹ So in general \vdash_{cw} seems a natural choice for answer set inference.

⁹ For example the main solvers such as `smodels`, `DLV` or `CLASP` do not implement a query answering mechanism.

¹⁰ In logical terms this constraint would be written $\alpha \rightarrow \perp$.

¹¹ Notice that for a non-ground, safe program an atomic query $q(a)$ is automatically false if a does not belong to the language of the program (even if q does), simply because grounding with the program constants is sufficient to generate all answer sets.

On the other hand, there are contexts where answer set semantics is used in a more open world setting, for example in the setting of hybrid knowledge bases [19] where non-monotonic rules are combined with ontologies formalised in description logics. For such systems a semantics in terms of equilibrium logic is provided in [5]. Here an entailment relation in the style of \vdash_{ow} might sometimes be more appropriate.

In general answer set semantics is defined only for coherent programs or theories. For finite, ground, coherent programs, by identifying \vdash_{AS} with \vdash_{cw} , we can apply Proposition 1 directly:

Corollary 2. *For coherent formulas α , (\vdash, \vdash) -interpolation in the form of Proposition 1 holds for entailment \vdash_{AS} in answer set semantics.*

5 Uniform Interpolation and Forgetting

A stronger form of interpolation known as *uniform* interpolation is also important for certain applications in computer science. As usual, given α, β with $\alpha \vdash \beta$, we are interested in interpolants γ such that

$$\alpha \vdash \gamma \ \& \ \gamma \vdash \beta \quad (10)$$

where $V(\gamma) \subseteq V(\alpha) \cap V(\beta)$. The difference now is that γ is said to be a *uniform interpolant* if (10) holds *for any* β in the same language such that $\alpha \vdash \beta$. A logic is said to have the uniform interpolation property if such uniform interpolants exist for all α, β .

While uniform interpolation fails in classical logic and in many non-classical logics, it may hold when certain restrictions are placed on the theory language where α is formulated and the query language containing β . For example it has been shown to hold for some description logics ([12]) where such syntactic restrictions apply. Even in ASP it turns out that a form of uniform interpolation holds for a very restricted query language, essentially one that allows just instance retrieval. We can show this by using some known results in ASP about the concept of *forgetting* [8] that is quite closely related to interpolation.

Variable forgetting, as studied in [8], is concerned with the following problem. Given a disjunctive logic program Π and a certain atom a occurring in Π , construct a new program, to be denoted by $\text{forget}(\Pi, a)$, that does not contain a but whose answer sets are in other respects as close as possible to those of Π . For the precise notion of closeness the reader is referred to [8], however some consequences will be evident shortly. In [8] the authors define $\text{forget}(\Pi, a)$ (as a generic term), show that such programs exist whenever Π is coherent, and provide different algorithms to compute such programs.

Given coherent Π and a in Π , the results $\text{forget}(\Pi, a)$, of forgetting about a in Π may be different but are always answer set equivalent. Moreover for our purposes they satisfy the following key property, where Π is coherent, a, b are distinct atoms in Π and as usual \vdash denotes nonmonotonic consequence,

$$\Pi \vdash b \Leftrightarrow \text{forget}(\Pi, a) \vdash b. \quad (11)$$

showing that indeed the answer sets of Π and $\mathbf{forget}(\Pi, a)$ are closely related.

To establish a version of uniform interpolation for the case of disjunctive programs and simple, atomic queries, we need to show that we can always find a $\Pi' = \mathbf{forget}(\Pi, a)$ such that $\Pi \sim \Pi'$. For this we can examine the first algorithm of [8] for computing $\mathbf{forget}(\Pi, a)$; this is also the simplest of the three algorithms presented. Let Π be a coherent program with rules of form (8) that we write as formulas of form (9) and let a be an atom in Π . The method for constructing a $\Pi' = \mathbf{forget}(\Pi, a)$ is as follows.

1. Compute the equilibrium models $E(\Pi)$.
2. Let E' be the result of removing a from each $\mathcal{M} \in E(\Pi)$.
3. Remove from E' any model that is non-minimal to form $E'' (= \{A_1, \dots, A_m\}$, say).
4. Construct a program Π' whose answer sets are precisely $\{A_1, \dots, A_m\}$ as follows:
 - for each A_i , set $\Pi_i = \{\neg \bar{A}_i \rightarrow a' : a' \in A_i\}$, where $\bar{A}_i = V(\Pi) \setminus A_i$.
 - Set $\Pi' = \Pi_1 \cup \dots \cup \Pi_m$.

We can now verify the desired property. Let \mathcal{L} be the simple query language composed of conjunctions of literals.

Proposition 9. *In equilibrium logic (or answer set programming) uniform interpolation holds for (coherent) disjunctive programs and queries in $\mathcal{L}(V(\Pi))$.*

Proof. To prove the claim we shall show the following. Let Π be a coherent disjunctive program and let $\mathcal{L} = \mathcal{L}(V)$ for some $V \subseteq V(\Pi)$. Then there is a program Π' such that $V(\Pi') = V$ and for any $\varphi \in \mathcal{L}$,

$$\Pi \sim \varphi \rightarrow (\Pi \sim \Pi' \ \& \ \Pi' \sim \varphi) \quad (12)$$

To begin, let Π and φ be as above with $\Pi \sim \varphi$. Let $X = \{a_1, \dots, a_n\} = V(\Pi) \setminus V$. Then we choose Π' to be the result of forgetting about X in Π , defined in [8] as follows:

$$\mathbf{forget}(\Pi, X) := \mathbf{forget}(\mathbf{forget}(\mathbf{forget}(\Pi, a_1), a_2), \dots, a_n),$$

and it is shown there that the order of the atoms in X does not matter. Now we know by (11) that for any atom $a \in V$ and any $i = 1, n$,

$$\Pi \sim a \Leftrightarrow \mathbf{forget}(\Pi, a_i) \sim a, \quad (13)$$

therefore

$$\Pi \sim a \rightarrow \mathbf{forget}(\Pi, X) \sim a. \quad (14)$$

Let Π' be $\mathbf{forget}(\Pi, X)$ as determined by algorithm 1 of [8] described above. It is easy to see by (13) and the semantics of \sim that (14) continues to hold where a is replaced by a negated atom $\neg b$ and therefore also by any conjunction of literals since a conjunction is entailed only if each element holds in every equilibrium model.¹² So

¹² As [8] points out, if an atom b is true in some answer set of $\mathbf{forget}(\Pi, a)$, then it must also be true in some answer set of Π , showing that (14) holds for literals.

it remains to show that $\Pi \sim \Pi'$. Again, it will suffice to show this entailment for one member of the sequence $\text{forget}(\Pi, a_i)$ and since the order is irrelevant wlog we can choose the first element $\text{forget}(\Pi, a_1)$ and show that $\Pi \sim \text{forget}(\Pi, a_1)$. We compute the programs Π_1, \dots, Π_m as in the algorithm. Then we need to check that $\Pi \sim \Pi_i$ for any $i = 1, \dots, n$, ie that for each $\mathcal{M} \in E(\Pi)$, $\mathcal{M} \models \{\neg \bar{A}_i \rightarrow a' : a' \in A_i\}$.

Consider $\mathcal{M} \in E(\Pi)$ where $\mathcal{M} = \langle T, T \rangle$. We distinguish two cases. (i) $A_i \subseteq T$. Then $\mathcal{M} \models a'$ for each $a' \in A_i$. It follows that $\mathcal{M} \models \neg \bar{A}_i \rightarrow a'$ for each $a' \in A_i$ and so $\mathcal{M} \models \{\neg \bar{A}_i \rightarrow a' : a' \in A_i\}$. Case (ii) $A_i \not\subseteq T$. Then T and A_i are incomparable. In particular we cannot have $T \subset A_i$ by the minimality property of A_i obtained in step 3. Hence $T \cap \bar{A}_i \neq \emptyset$. Choose $a'' \in T \cap \bar{A}_i$. Then $\mathcal{M} \models a''$, so $\mathcal{M} \not\models \neg a''$ and hence $\mathcal{M} \not\models \neg \bar{A}_i$. Consequently, for any a' , $\mathcal{M} \models \neg \bar{A}_i \rightarrow a'$ and so $\mathcal{M} \models \{\neg \bar{A}_i \rightarrow a' : a' \in A_i\}$.

It follows that for any i , $\Pi \sim \Pi_i$ and so by construction $\Pi \sim \Pi'$, which establishes the proposition. \square

5.1 Extending the query language

If we establish uniform interpolation in ASP using the method of forgetting, as defined in [8], it seems clear that we cannot extend in a non-trivial way the expressive power of the query language \mathcal{L} . Since the method of forgetting a in Π removes non-minimal sets from $E(\Pi)$ (once a has been removed), an atom b might be true in some equilibrium model of Π but not in any equilibrium model of $\text{forget}(\Pi, a)$. Hence we might have a disjunction, say $a \vee b$, derivable from Π but not from $\text{forget}(\Pi, a)$.

On the other hand, the property of uniform interpolation certainly holds for any $\mathcal{L}(V)$ even without the condition that $V \subseteq V(\Pi)$. Suppose that $\Pi \sim \varphi$ where $V(\varphi) \setminus V(\Pi) \neq \emptyset$, say $V(\varphi) \setminus V(\Pi) = \{b_1, \dots, b_k\}$. Then b_1, \dots, b_k are false in all equilibrium models of Π . Trivially, if b is not in $V(\Pi)$ we can regard the result of forgetting about b in Π as just Π . So we can repeat the proof of Proposition 9, but now setting $X = \{V(\Pi) \setminus V\} \cup \{V \setminus V(\Pi)\}$. All the relevant properties will continue to hold.

An interesting open question is whether we can extend the theory language to include more general kinds of program rules allowing negation in the head. Accommodating these kinds of formulas would constitute an important generalisation since they amount to a normal form in equilibrium logic. However, the answer sets of such programs do not satisfy the minimality property that holds for the answer sets of disjunctive programs, so it is clear that the definition of forgetting would need to be appropriately modified - a task that we do not attempt here.

6 Literature and Related Work

The interpolation theorem for classical logic is due to Craig [6]; it was extended to intuitionistic logic by Schütte [20]. Maksimova [16] characterised the super-intuitionistic propositional logics possessing interpolation. A modern, comprehensive treatment of interpolation in modal and intuitionistic logics can be found in the monograph [10] by Gabbay and Maksimova.

In non-monotonic logics, interpolation has received little attention. A notable exception is an article [1] by Amir establishing some interpolation properties for circumscription and default logic. By the well-known relation between the answer sets of disjunctive programs and the extensions of corresponding default theories, he also derives a form of interpolation for ASP. With regard to answer set semantics, the approach of [1] is quite different from ours. Since it is founded on an analysis of default logic, it uses classical logic as an underlying base. So Amir's version of interpolation is a form of (4) where L is classical logic; there is no requirement that \vdash_L form a well-behaved sublogic of \vdash , eg a deductive base. As Amir remarks, one cannot deduce in general from property (5) that $\alpha \vdash \beta$. However if L is classical logic one cannot even deduce $\alpha \vdash \beta$ from (4). More generally, there is no counterpart to our Proposition 1 in this case. Another difference with respect to our approach is that [1] does not discuss the nature of the \vdash relation for ASP in detail, in particular how to understand $\Pi \vdash \varphi$ in case φ contains atoms not present in the program Π . In fact, if we interpret \vdash_{AS} as in section 5 above, it is easy to refute (\vdash, \vdash_L) -interpolation where L is classical logic. Let Π be the program $B \leftarrow \neg A$ and q the query $B \wedge \neg C$. Then clearly $\Pi \vdash_{AS} q$, but there is no formula in the vocabulary B that would classically entail $\neg C$. Under any interpretation of answer set inference such that atoms not in the program are regarded as false, (\vdash, \vdash_L) -interpolation would be refuted.

7 Conclusions

We have discussed two kinds of interpolation properties for non-monotonic inference relations and shown that these properties hold in turn for the two different inference relations that we can associate with propositional equilibrium logic. In each case we use the fact that the collection of equilibrium models is definable in the logic **HT** of here-and-there and that this logic possesses the usual form of interpolation. One of the forms of inference studied seems to be in many cases an appropriate concept to associate with answer set programming, although in general ASP systems are not tailored to query answering or deduction. Using results from [8] about variable forgetting in ASP, we could also show how the property of uniform interpolation holds for disjunctive programs and a restricted query language.

A forthcoming, extended version of this paper will deal with interpolation in quantified equilibrium logic and non-ground ASP. Results similar to the propositional case can be established, providing that the class of equilibrium models is (first-order) definable, as for instance in the case of safe theories.

References

1. E. Amir. Interpolation Theorems for Nonmonotonic Reasoning Systems. In S. Flesca et al. (Eds), JELIA 2002, LNAI 2424, pp. 233–244, Springer, 2002.
2. M. Balduccini, M. Gelfond, M. Nogueira. A-Prolog as a tool for declarative programming. Proc. of SEKE 2000.
3. C. Baral. Knowledge Representation, Reasoning and Declarative Problem Solving. CUP, 2002.

4. J. Bicarregui, T. Dimitrakos, D. Gabbay & T. Maibaum. Interpolation in Practical Formal Development. *Logic J. of the IGP*, 9, no. 2 (2001), 247-259.
5. J. de Bruijn, D. Pearce, A. Polleres & A. Valverde. Quantified Equilibrium Logic and Hybrid Rules. In M. Marchiori et al (eds), *Proc. of RR2007*, Springer LNCS 4524, 2007, 58-72.
6. W. Craig. Linear reasoning. A new form of the Herbrand-Gentzen theorem. *J. Symb. Logic*, 22 (1957), 250-268.
7. J. Dietrich. Deductive Bases of Nonmonotonic Inference Operations. NTZ Report, University of Leipzig, 1994.
8. T. Eiter & K. Wang. Semantic forgetting in answer set programming. *Art. Int.* 172 (2008), 1644-1672.
9. P. Ferraris. Answer Sets for Propositional Theories. In C. Baral et al (eds) *Proc. of LPNMR 05*, Springer LNAI 3662, 2005.
10. D. Gabbay and L. Maksimova. *Interpolation and Definability: Modal and Intuitionistic Logics*. Oxford Logic Guides: 46, OUP, Oxford, 2005
11. M. Gelfond & V. Lifschitz. The Stable Model Semantics for Logic Programs. In K. Bowen & R. Kowalski (eds), *Proc. of ICLP 88*, MIT Press, 1070-1080.
12. R. Kontchakov, F. Wolter, & M Zakharyashev. A logic-based framework for ontology comparison and module extraction in DL-Lite *J. Art. Int.*, to appear.
13. V. Lifschitz, D. Pearce & A. Valverde. Strongly equivalent logic programs. *ACM Trans. on Comp. Logic*, 2:526-541, 2001.
14. V. Lifschitz, L. Tang & H. Turner. Nested Expressions in Logic Programs. *Ann. Math. & Art. Int.*, 25(3-4): 369-389, 1999.
15. D. Makinson. General Patterns in Nonmonotonic Reasoning. In D. Gabbay et al (eds), *Handbook of Logic in Artificial Intelligence*, Clarendon Press, Oxford, 1994.
16. L. Maksimova. Craig's Interpolation theorem and amalgamated varieties. *Doklady AN SSR*, 237:1281-1284, 1977.
17. D. Pearce. A New Logical Characterisation of Stable Models and Answer Sets. In J. Dix et al (eds), *Proc. of NMELP 96*, Springer, LNAI 1216, 1997, 57-70.
18. D. Pearce. Equilibrium logic. *Ann. Math. & Art. Int.*, 47 (2006), 3-41.
19. R. Rosati. Semantic and Computational Advantages of the Safe Integration of Ontologies and Rules. *Proc of PPSWR 2005*.
20. K. Schütte. Der Interpolationsatz der Intuitionistischen Prädikatenlogik. *Math. Ann.* 148 (1962), 192-200.
21. D. van Dalen. *Logic and Structure*, 4th Ed. Springer, 2004.